

Genome-wide association study of 107 phenotypes in a common set of *Arabidopsis thaliana* inbred lines

Susanna Atwell^{**}, Yu S. Huang^{1,*}, Bjarni J. Vilhjálmsson^{1,*}, Glenda Willems^{1,*}, Matthew Horton², Yan Li², Dazhe Meng¹, Alexander Platt¹, Aaron M. Tarone¹, Tina T. Hu¹, Rong Jiang¹, N. Wayan Muliyati², Xu Zhang², Muhammad Ali Amer¹, Ivan Baxter³, Benjamin Brachi⁴, Joanne Chory^{5,6}, Caroline Dean⁷, Marilyne Debieu⁸, Juliette de Meaux⁸, Joseph R. Ecker⁶, Nathalie Faure⁴, Joel M. Kniskern², Jonathan D. G. Jones⁹, Todd Michael⁶, Adnane Nemri⁹, Fabrice Roux^{2,4}, David E. Salt¹⁰, Chunlao Tang¹, Marco Todesco¹¹, M. Brian Traw², Detlef Weigel¹¹, Paul Marjoram¹², Justin O. Borevitz², Joy Bergelson², Magnus Nordborg^{1,13}

December 8, 2009

Although pioneered by human geneticists as a potential solution to the challenging problem of finding the genetic basis of common human diseases^{1,2}, advances in genotyping and sequencing technology have made genome-wide association (GWA) studies an obvious general approach for studying the genetics of natural variation and traits of agricultural importance. They are particularly useful when inbred lines are available because once these lines have been genotyped, they can be phenotyped multiple times, making it possible (as well as extremely cost-effective) to study many different traits in many different environments, while replicating the phenotypic measurements to reduce environmental noise. Here we demonstrate the power of this approach by carrying out a GWA study of 107 phenotypes in *Arabidopsis thaliana*, a widely distributed, predominantly selfing model plant, known to harbor considerable genetic variation for many adaptively important traits³. Our results are dramatically different from those of human GWA studies in that we identify many common alleles with major effect, but they are also, in many cases, harder to interpret because confounding by complex genetics and population structure make it difficult to distinguish true from false associations. However, *a priori* candidates are significantly overrepresented among these associations as well, making many of them excellent candidates for follow-up experiments by the *Arabidopsis* community. Our study clearly demonstrates the feasibility of GWA studies in *A. thaliana*, and suggests that the approach will be appropriate for many other organisms.

The genotyped sample (Supplementary Table 1) includes a core set of 95 lines⁴ for which a wide variety of phenotypes were available, plus a second set of 96 for which many phenotypes related to flowering were available⁵. The lines were genotyped using a custom Affymetrix SNP-chip containing 250,000 SNPs⁶. Since the genome of *A. thaliana* is around 120 million bp and

^{*}Molecular and Computational Biology, University of Southern California, Los Angeles, California, USA. ²Department of Ecology & Evolution, University of Chicago, Chicago, Illinois, USA. ³Bindley Bioscience Center, Purdue University, West Lafayette, Indiana, USA. ⁴Laboratoire de Génétique et Evolution des Populations Végétales, UMR CNRS 8016, Université des Sciences et Technologies de Lille 1, Villeneuve d'Ascq Cedex, France. ⁵Howard Hughes Medical Institute, La Jolla, California, USA. ⁶Plant Biology Laboratory, The Salk Institute for Biological Studies, La Jolla, California, USA. ⁷Department of Cell and Development Biology, John Innes Centre, Norwich, UK. ⁸Max Planck Institute for Plant Breeding Research, Cologne, Germany. ⁹Sainsbury Laboratory, Norwich, UK. ¹⁰Purdue University, West Lafayette, Indiana, USA. ¹¹Department of Molecular Biology, Max Planck Institute for Developmental Biology, Tübingen, Germany. ¹²Department of Preventive Medicine, Keck School of Medicine, University of Southern California, Los Angeles, California, USA. ¹³Gregor Mendel Institute, Vienna, Austria.

^{*}These authors contributed equally to the paper.

the extent of LD comparable to that in humans^{7,8}, the resulting SNP density of one SNP per 500 bp is considerably higher than is commonly used in human studies⁶.

To evaluate the feasibility of GWA studies in this organism, a variety of phenotypes were generated or assembled. The phenotypes broadly fell into four categories: 23 were related to flowering under different environmental conditions; 23 were related to defense, ranging from recognition of specific bacterial strains to trichome density; 18 were element concentrations measured using inductively coupled plasma mass spectroscopy (“ionomics”); and 43 were loosely defined developmental traits, including dormancy and plant senescence. For details about each phenotype, see Supplementary Tables 2–5. The flowering phenotypes are generally strongly positively correlated, and are also negatively correlated with some other phenotypes, e.g., those related to size at flowering (see Supplementary Fig. 9).

We first assessed evidence of association between each SNP and phenotype using the non-parametric Wilcoxon Test (Fisher’s Exact Test was used for the small number of phenotypes that were categorical rather than quantitative). A major difference between our study and the human GWA studies published to date is that our study population is heavily structured, and there is thus every reason to expect elevated false-positive rates⁹. Indeed, most phenotypes gave rise to a distribution of p-values that was strongly skewed toward zero (Supplementary Section 2.1.6). Fig. 1a shows the number of distinct peaks of association identified for each phenotype using different p-value thresholds, as well as the number expected by chance alone. There is an excess of strong associations across phenotypes, as expected given the presence of confounding population structure (although we also expect some of these associations to be true). Furthermore, the degree of confounding varies greatly between phenotypes. Phenotypes related to flowering are generally more strongly affected, as would be expected given the correlation between flowering and geographic origins.

The population structure in our sample is highly complex, involving patterns of relatedness on all scales (see Supplementary Fig. 4). As in previous studies⁹, we found that, at least in terms of producing a p-value distribution that does not show obvious signs of confounding, statistical methods commonly used to control for population structure in human genetics^{10,11} fail to correct for population structure, whereas the mixed-model approach introduced by maize geneticists¹² appears to perform well (see Supplementary Section 2.1.6). Fig. 1b shows the number of peaks of association identified using this approach (as implemented in the program EMMA¹³). The excess of nominally significant association for flowering-related phenotypes has been eliminated, as would be expected if this excess were mostly due to confounding by population structure. There is a marked reduction in the number of associations of moderate significance (e.g., around 10^{-4}) across phenotypes, but the excess of highly significant associations clearly persists (or has even become greater). This is precisely what would be expected from an increase in statistical power by switching to a parametric method that reduces confounding. It is tempting to conclude that most of these extreme p-values must represent true associations, but there are reasons to be skeptical. First, although EMMA appears to produce a p-value distribution that conforms to the null-expectation (except for extreme values: see Supplementary Figs. 12–118), it seems almost certain that some confounding remains (see below). Second, simulation studies suggest that the p-values produced by EMMA are not always well estimated and should be interpreted with caution (see Supplementary Section 2.1.5).

It is thus not straightforward to distinguish true from spurious associations, regardless of whether we correct for population structure. There is no doubt, however, that there is real signal in the data. Indeed, regardless of method used, six phenotypes yield single, strong peaks of association that are obvious by eye. In all cases, the association results effectively identify single genes, and correspond to known functional polymorphisms. An example of this is shown in Fig. 2, where the hypersensitive response to the bacterial avirulence gene *AvrRpm1* directly identifies the corresponding resistance gene *RESISTANCE TO P. SYRINGAE PV MACULICOLA 1 (RPM1)*¹⁴. Similar results were obtained for other disease resistance (*R*) responses, sodium concentration, lesioning, and *FRIGIDA (FRI)* expression (see Supplementary Figs. 35, 37, 60, 76, and 21, respectively).

More generally, SNPs closely linked to genes that are *a priori* likely to be responsible for a particular phenotype are significantly over-represented among SNPs associated with that phenotype.

For many of the phenotypes analyzed it is possible to predict which genes might be important in natural variation based on existing functional knowledge. For these phenotypes we determined which of our SNPs were located within 20 kb of an *a priori* candidate, and tested whether these SNPs were over-represented among nominally significant associations. Fig. 3 illustrates the procedure for flowering time at 10°C (FT10). For example, SNPs with a p-value less than 10^{-3} from EMMA and a p-value less than 10^{-5} from the Wilcoxon test are 2.7 times more likely to be close to candidate genes than are randomly chosen SNPs. This simultaneously demonstrates that background functional knowledge about flowering pathways helps predict which genes are involved in natural variation and that our GWA results identify many true associations. Indeed, by assuming that all associations not involving *a priori* candidates are false, we see that the inverse of the enrichment ratio provides a crude upper bound for the false-positive rate among the *a priori* candidates (see Supplementary Section 3.2). Continuing the example in Fig. 3, no more than 40% of the candidate SNPs that are significant using both tests are false. This is an upper bound because it seems almost certain that many of the (much larger set of) strong associations that are not close to *a priori* candidates will also turn out to be real.

As illustrated in Fig. 3, *a priori* candidates are over-represented among strongly associated SNPs regardless of whether we correct for population structure or not, and the SNPs identified are not necessarily the same. Enrichment is clearly greatest among SNPs that are strongly associated using both methods, however. This is true across the flowering-related phenotypes (Supplementary Fig. 10), and it is thus clear that both methods have utility. More stringent thresholds typically yield stronger enrichment, but the variance also increases because the number of significant genes decreases, and there is thus no simple relationship between degree of enrichment and its statistical significance (Supplementary Fig. 11). Results for the other phenotypes are consistent with those for the flowering-related traits, but the candidate gene lists are too short for statistical analysis (Supplementary Section 3.1).

An additional problem in identifying true positives was the existence of complex peaks of association. While many peaks were sharply defined and clearly identified a small number of genes (illustrated in Fig. 2b), others were much more diffuse, sometimes covering several hundred kb without a clear center. Fig. 4 shows an example of such a peak, and also suggests an explanation for their existence. The figure shows the pattern of association with *FLOWERING LOCUS C* (*FLC*) expression in the chromosomal region containing the vernalization-response gene *FRI*. Polymorphisms in *FRI* are known to affect flowering time partly through their effect on expression of *FLC*^{5,15}. SNPs in the *FRI* region should thus be associated with *FLC* expression. This is indeed the case, but rather than a single peak of association centered on *FRI*, we have a mountain range covering 500 kb and on the order of a hundred genes (Fig. 4a).

That *FRI* should be surrounded by a wide peak of association is, in itself, not surprising given that the two common loss-of-function alleles at *FRI* appear to have been the subject of recent positive selection¹⁶. Indeed, the entire range collapses if these two alleles are added as co-factors to the model (Fig. 4d). More surprising is the fact that these two causal polymorphisms do not have the strongest association within the region. If we reduce allelic heterogeneity by factoring out one or the other of the two alleles, the significance of the remaining polymorphism increases, but it is still not the most significant in the region (Fig. 4b–c). A likely explanation for this is that some SNPs in the region are positively correlated (in linkage disequilibrium) with one of the *FRI* alleles (because of linkage) and the genomic background (because of population structure). This dual confounding is sufficient to make some of these SNPs more strongly associated with the phenotype than the true positives.

Given the difficulties described above, deciding which associations are worth following-up must necessarily be highly subjective. The strongest associations do not always correspond to obvious candidates and are perhaps more interesting than associations in genes with known function. However, in the absence of further evidence there is little point in discussing these associations. Supplementary Table 6 lists some of the most promising associations: additional *a posteriori* candidates for each phenotype are given in Supplementary Figs. 12–118. The genes listed were selected based on annotation from within a 20 kb window surrounding each of the top 500 most strongly associated SNPs (with minor allele frequency ≥ 0.1 for EMMA), distinguishing between those that had been considered candidates *a priori* and those that had not (the latter category is

marked with asterisks in the tables). As demonstrated in Fig. 3, we expect a high fraction of the associated *a priori* candidates to be real. The full data are available through the project website (<http://arabidopsis.usc.edu>).

For the flowering-related phenotypes, one of the most striking findings was the strong correlation between phenotypes generated under very different growth-chamber and greenhouse conditions (Supplementary Table 2 and Supplementary Fig. 9: note that phenotypes from a field experiment were much less strongly correlated). As expected given the correlation in phenotypes, there are several regions of association that are shared across the majority of the flowering phenotypes (Supplementary Fig. 119). These regions vary considerably in width, and many of them are complex in the sense of Fig. 4, perhaps as consequence of strong selection. As expected given the results presented in Fig. 3, several of these regions coincided with *a priori* candidates, like *FRI*¹⁵ and *FLC*¹⁷ (see Supplementary Table 6). Another interesting candidate is *DELAY OF GERMINATION 1 (DOG1)*¹⁸, which, though not originally thought to be involved with flowering, is highly associated with 20 different flowering phenotypes.

Among the other phenotypes, three previously identified resistance (*R*) genes polymorphisms were readily identified⁸, as were genes known to be involved with variation in sodium (Na)¹⁹ and molybdenum (Mo)²⁰ levels (Supplementary Table 6). Four genes known to be involved trichome-formation were strongly associated with both trichome phenotypes: one of these associations has recently been confirmed²¹. Finally, *ACD6*, which has been experimentally shown to be directly involved in lesioning²², is detected here as associated with several lesioning and chlorosis phenotypes. This association has also recently been experimentally confirmed²³.

By the standards of human GWA studies, the sample sizes used in this study (~96 or ~192 lines) are tiny. It may thus seem surprising that we are able to map anything at all. However, power depends on the genetic architecture of the traits as well, and this works in our favor in at least two ways. First, we clearly find common alleles of major effect. Although effect sizes are hard to estimate for the same reason p-values are, we note, for example, that in 44 phenotypes at least one of the 50 most strongly associated SNPs with a minor allele frequency greater than 15% explain more than 20% of the phenotypic variance (effect-size estimates and allele frequencies for every association are on the project website). This is very different from human studies, which have generally identified only polymorphisms of very small phenotypic effect. The difference is likely due to the fact that, whereas human studies have focused on traits that are either deleterious or under strong stabilizing selection (for which a very strong trade-off between allelic effect and frequency is expected²⁴), we are working with adaptively important traits. Indeed, human GWA studies focusing on traits like skin color seem to yield results more like those presented here^{25,26}.

Second, our study takes full advantage of the fact that we are working with inbred lines that can be grown in replicate under controlled conditions, making it possible to study multiple phenotypes while controlling environmental noise. Partially as a result of this, heritabilities for the traits studied are generally high, ranging from 42% for aphid number to over 99% for several flowering traits (Supplementary Table 7).

Without these advantages, the amount of genotyping required would have made a study like the present one prohibitively expensive. That said, there is little doubt that power in our study is severely limited by sample size. Simulation studies indicate that our power to detect alleles similar to those actually detected in the study is often no more than 30–40% using a sample size of 96 (results not shown). Increasing the sample size to 192 typically more than doubles power. Well over 1,000 lines genotyped with our 250k SNP chip will soon be available: we look forward to seeing associations from these lines. Efforts are also underway to sequence the genomes of all these lines (<http://www.1001genomes.org>) — this will greatly facilitate follow-up studies, but we do not expect a massive increase in power as a result of this, because the SNP density used here seems adequate.

Power also depends on sample composition. In comparison with typical human GWA studies, our sample is characterized by extremely strong population structure. This is expected given that our global sample was collected partly to study population structure in *A. thaliana*⁴. GWA studies that utilize different samples (including regional, more homogeneous ones) are under way.

The fact that population structure can cause confounding and lead to an elevated false-positive rate is well known and the relative advantages of alternative statistical methods to correct for this

have been much debated^{10,11,12,27}. We feel that the discussion of this phenomenon has often been misleading, in that population structure is neither necessary nor sufficient for confounding to occur. At least for complex traits, the problem is better thought of as model mis-specification: when we carry out GWA analysis using a single SNP at a time (as is done in this, and most other GWA studies to date), we are in effect modeling a multi-factorial trait as if it were due to a single locus. The polygenic background of the trait is ignored, as are other unobserved variables. This kind of marginal analysis causes no problem as long as the background is adequately captured by a variance term (or similar), but if the background variables are correlated with the SNP included in the model, bias will result. Population structure will lead to correlations (i.e., linkage disequilibrium) between unlinked loci, and this will usually (but not always²⁸) lead to confounding. Positive correlations are also expected as a result of strong selection. Both factors are likely to be important in the present study: for example, it is easy to imagine that plants from northern Sweden will tend to share cold-adaptive alleles at many causal loci as a result of selection, and marker alleles genome-wide as a result of demographic history.

This way of thinking about the problem helps us interpret many of the results presented above. First, it becomes clear why GWA works so well for traits that are monogenic, or at least are mostly due to a single major locus. Examples in our study include the *R* gene responses (*RPM1*, *RPS2*, and *RPS5*), *FRI* expression (*FRI* itself), and lesioning (*ACD6*). In all cases, GWA yields unambiguous results regardless of whether we correct for population structure. The reason is not that there is no confounding in these cases. The problem that has received so much attention in human genetics — inflated significance among unlinked, non-causal loci — is clearly present, but with truly genome-wide coverage this is not very important, because the true positive is expected to show the strongest association.

Second, it helps us explain the occurrence of broad, complex regions of association. As exemplified in Fig. 4, these can arise when SNPs in a region containing a major causative allele are positively correlated not only with that causative allele (due to linkage disequilibrium in the narrow sense), but also with the genomic background (because of population structure and/or natural selection). The paradoxical consequence is that instead of a single peak of association centered on the causative locus, we expect a complex mountain landscape where many non-causal markers will show stronger association than the causative allele itself. This makes it difficult to identify the causal variant within such regions. This type of confounding does not appear to have been recognized in the literature, and probably deserves more attention.

Third, it is clear that we should not expect statistical methods that are designed to take genome-wide patterns of relatedness into account^{10,11,12} to correct for confounding that is due to selection generating correlations between causal loci. These types of methods will work only to the extent that the loci responsible for the genomic background have allele-frequency distributions that are similar to those of non-causal loci, which is expected only if selection on each locus is weak. Accurate estimation of the size of the effects of the many candidate polymorphisms identified here (including distinguishing it from zero) will require either crosses or transgenic experiments. Any cross will eliminate long-range linkage disequilibrium, and short-range linkage disequilibrium can be overcome by choosing appropriate parental strains. For example, the *FRI* region in Fig. 4 also contains a promising association at *CRYPTIC PRECOCIOUS* (*CRP*), less than 100 kb away from *FRI* (see Supplementary Fig. 127). If polymorphism at *FRI* is taken into account, *CRP* is no longer significantly associated, but this does not necessarily mean that the association is spurious. The two genes are too closely linked to be separated using standard crosses, but we can select parental strains that segregate for *CRP* but not *FRI*.

Such crosses are currently being carried out, by us and by other members of the *Arabidopsis* community. We also anticipate that many more phenotypes will be generated and added to our public database. By combining results from GWA, linkage mapping, and perhaps also intermediate phenotypes (e.g., expression data), it will be possible to make progress on deconstructing the regulatory networks that determine natural variation.

As genotyping and sequencing costs continue to decrease, GWA studies will become a standard tool for dissecting natural variation. It is thus important to recognize their limitations. The problems raised here are not unique to *A. thaliana*. GWA alone will often not allow accurate estimate of allelic effects. It must also be remembered that all mapping studies are biased in the sense

that they can only detect alleles that explain a sufficient fraction of the variation in the mapping population²⁹. The present study can only detect alleles that are reasonably common in our global sample. A GWA study using a more local sample would undoubtedly uncover more variants that are locally common, and linkage mapping will identify major polymorphisms that happen to be segregating in the cross, even if one of the alleles is extremely rare in natural populations. The “genetic architecture” of a trait depends on the population studied. In order to determine how genetic architecture affects selection and evolution, we thus also need to understand the spatial and temporal scales over which selection is important.

Methods summary

Because slightly different sets were used in different phenotyping experiments, the total number of lines used was 199 (Supplementary Table 1). Genotyping was done using standard protocols, and a combination of SNP calling and imputation algorithms were used to analyze the results (see Supplementary Section 1). We called 216,130 SNPs, at an estimated error rate of 1.6%. GWA analysis was done with and without correction for confounding. For the former, a mixed-model¹² implemented in the program EMMA¹³ was used. For the latter, Wilcoxon’s test was used for ordinal data, and Fisher’s Exact Test for categorical data. Enrichment for candidate genes was investigated using lists of *a priori* candidates identified from the literature.

Full Methods and any associated references are part of the Supplementary Information available online at www.nature.com/nature.

REFERENCES

1. Hirschhorn, J. N. and Daly, M. J. *Nat Rev Genet* **6**, 95–108 Feb (2005).
2. Wellcome Trust Case-Control Consortium. *Nature* **447**, 661–78 Jun (2007).
3. Koornneef, M., Alonso-Blanco, C., and Vreugdenhil, D. *Annu. Rev. Plant Biol.* **55**, 141–172 (2004).
4. Nordborg, M., Hu, T. T., Ishino, Y., Jhaveri, J., Toomajian, C., Zheng, H., Bakker, E., Calabrese, P., Gladstone, J., Goyal, R., Jakobsson, M., Kim, S., Morozov, Y., Padhukasahasram, B., Plagnol, V., Rosenberg, N. A., Shah, C., Wall, J., Wang, J., Zhao, K., Kalbfleisch, T., Schultz, V., Kreitman, M., and Bergelson, J. *PLoS Biol.* **3**, e196 (2005).
5. Shindo, C., Aranzana, M. J., Lister, C., Baxter, C., Nicholls, C., Nordborg, M., and Dean, C. *Plant Physiol.* **138**, 1163–1173 (2005).
6. Kim, S., Plagnol, V., Hu, T. T., Toomajian, C., Clark, R. M., Ossowski, S., Ecker, J. R., Weigel, D., and Nordborg, M. *Nature Genet.* **39**, 1151–1155 (2007).
7. Nordborg, M., Borevitz, J. O., Bergelson, J., Berry, C. C., Chory, J., Hagenblad, J., Kreitman, M., Maloof, J. N., Noyes, T., Oefner, P. J., Stahl, E. A., and Weigel, D. *Nature Genet.* **30**, 190–193 (2002).
8. Aranzana, M. J., Kim, S., Zhao, K., Bakker, E., Horton, M., Jakob, K., Lister, C., Molitor, J., Shindo, C., Tang, C., Toomajian, C., Traw, B., Zheng, H., Bergelson, J., Dean, C., Marjoram, P., and Nordborg, M. *PLoS Genet.* **1**, e60 (2005).
9. Zhao, K., Aranzana, M. J., Kim, S., Lister, C., Shindo, C., Tang, C., Toomajian, C., Zheng, H., Dean, C., Marjoram, P., and Nordborg, M. *PLoS Genet.* **3**, e4 (2007).
10. Pritchard, J. K., Stephens, M., Rosenberg, N. A., and Donnelly, P. *Am J Hum Genet* **67**, 170–181 (2000).

11. Price, A. L., Patterson, N. J., Plenge, R. M., Weinblatt, M. E., Shadick, N. A., and Reich, D. *Nat Genet* **38**, 904–9 (2006).
12. Yu, J., Pressoir, G., Briggs, W., Vroh Bi, I., Yamasaki, M., Doebley, J., McMullen, M., Gaut, B., Nielsen, D., Holland, J., Kresovich, S., and Buckler, E. *Nat Genet* **38**, 203–208 (2006).
13. Kang, H. M., Zaitlen, N. A., Wade, C. M., Kirby, A., Heckerman, D., Daly, M. J., and Eskin, E. *Genetics* **178**, 1709–23 Mar (2008).
14. Grant, M. R., Godiard, L., Straube, E., Ashfield, T., Lewald, J., Sattler, A., Innes, R. W., and Dangl, J. L. *Science* **269**, 843–846 (1995).
15. Johanson, U., West, J., Lister, C., Michaels, S., Amasino, R., and Dean, C. *Science* **290**, 344–347 (2000).
16. Toomajian, C., Hu, T. T., Aranzana, M. J., Lister, C., Tang, C., Zheng, H., Calabrese, P., Dean, C., and Nordborg, M. *PLoS Biol.* **4**, e137 (2006).
17. Michaels, S. D. and Amasino, R. M. *Plant Cell* **11**, 949–956 (1999).
18. Bentsink, L., Jowett, J., Hanhart, C. J., and Koornneef, M. *Proc. Natl. Acad. Sci. USA* **103**(45), 17042–17047 (2006).
19. Rus, A., Baxter, I., Muthukumar, B., Gustin, J., Lahner, B., Yakubova, E., and Salt, D. E. *Plos Genetics* **2**, 1964–1973 (2006).
20. Baxter, I., Muthukumar, B., Park, H. C., Buchner, P., Lahner, B., Danku, J., Zhao, K., Lee, J., Hawkesford, M. J., Guerinot, M. L., and Salt, D. E. *Plos Genetics* **4** (2008).
21. Hilscher, J., Schlötterer, C., and Hauser, M.-T. *Current Biology* (2009). To appear.
22. Lu, H., Rate, D. N., Song, J. T., and Greenberg, J. T. *Plant Cell* **15**, 2408–2420 (2003).
23. Todesco, M., Balasubramanian, S., Hu, T. T., Epple, P., Kuhns, C., Sureshkumar, S., Schwartz, C., Traw, M. B., Lanz, C., Laitinen, R. A. E., Chory, J., Lipka, V., Borevitz, J. O., Bergelson, J., Dangl, J. L., Nordborg, M., and Weigel, D. Under review for *Nature*, (2009).
24. Manolio, T. A., Collins, F. S., Cox, N. J., Goldstein, D. B., Hindorff, L. A., Hunter, D. J., McCarthy, M. I., Ramos, E. M., Cardon, L. R., Chakravarti, A., Cho, J. H., Guttmacher, A. E., Kong, A., Kruglyak, L., Mardis, E., Rotimi, C. N., Slatkin, M., Valle, D., Whittemore, A. S., Boehnke, M., Clark, A. G., Eichler, E. E., Gibson, G., Haines, J. L., Mackay, T. F. C., Mccarroll, S. A., and Visscher, P. M. *Nature* **461**, 747–53 (2009).
25. Han, J., Kraft, P., Nan, H., Guo, Q., Chen, C., Qureshi, A., Hankinson, S. E., Hu, F. B., Duffy, D. L., Zhao, Z. Z., Martin, N. G., Montgomery, G. W., Hayward, N. K., Thomas, G., Hoover, R. N., Chanock, S., and Hunter, D. J. *PLoS Genet* **4**(5), e1000074 May (2008).
26. Stokowski, R. P., Pant, P. V. K., Dadd, T., Fereday, A., Hinds, D. A., Jarman, C., Filsell, W., Ginger, R. S., Green, M. R., van der Ouderaa, F. J., and Cox, D. R. *Am J Hum Genet* **81**(6), 1119–32 Dec (2007).
27. Devlin, B. and Roeder, K. *Biometrics* **55**, 997–1004 (1999).
28. Rosenberg, N. and Nordborg, M. *Genetics* **173**, 1665–1678 (2006).
29. Nordborg, M. and Weigel, D. *Nature* **456**, 720–723 (2008).

Acknowledgments We thank B. Carvalho for his advice on how to modify the Oligo package. This work was primarily supported by NSF DEB-0519961 (J.B., M.N.), NIH GM073822 (J.O.B.), and NSF DEB-0723935 (M.N.). Additional support was provided by the Dropkin foundation, NIH GM057994 and NSF MCB-0603515 (J.B.), the Max Planck Society (D.W., M.T.), the Austrian Academy of Sciences (M.N.), the University of Lille 1 (F.R.), NIH GM078536 and NIH P42ES007373 (D.E.S.), NIH GM62932 (J.C. and D.W.), the Howard Hughes Medical Institute (J.C.), DFG SFB 680 (J.d.M.), a Marie-Curie International Outgoing Fellowship “ANAVACO” (Project Nr. 220833; G.W.), and a Gottfried Wilhelm Leibniz Award of the DFG (D.W.). The project would not have been possible without the existence of TAIR (<http://arabidopsis.org>).

Author Contributions J.R.E. and D.W. generated the SNPs used by this project. S.A., M.H., Y.L., N.W.M., X.Z., J.O.B., and J.B. were responsible for the experimental aspects of genotyping. Y.S.H., B.J.V., M.H., T.T.H., R.J., X.Z., M.A.A., P.M., J.O.B., J.B., and M.N. were responsible for data management and the bioinformatics pipeline. S.A., I.B., B.B., J.C., C.D., M.D., J.d.M., N.F., J.M.K., J.D.G.J., T.M., A.N., F.R., D.E.S., C.T., M.T., M.B.T., D.W., J.B., and M.N. were responsible for phenotyping. S.A., Y.S.H., B.J.V., G.W., D.M., A.P., A.M.T., P.M., and M.N. carried out the GWA analyses. Y.S.H. and D.M. developed the project website. M.N. wrote the paper with major contributions from S.A., Y.S.H., B.J.V., G.W., A.P., and J.B.. J.O.B., J.B., and M.N. conceived and supervised the project.

Author Information Reprints and permissions information is available at www.nature.com/reprints. Correspondence and requests for materials should be addressed to M.N. (magnus.nordborg@gmi.oeaw.ac.at).

Figure legends

Figure 1 | The number of associations identified using different p-value thresholds for each phenotype. For each phenotype, the numbers of distinct peaks of association significant at nominal p-value thresholds of 10^{-4} , 10^{-5} , \dots , 10^{-9} are shown. The number of SNPs (out of 250,000) that would be expected to exceed each threshold is shown for comparison. **a**, No correction for population structure (non-parametric Wilcoxon Test). **b**, Correction for population structure (parametric mixed model [EMMA]).

Figure 2 | GWA analysis of hypersensitive response to the bacterial elicitor *AvrRpm1*. **a**, Genome-wide p-values from Fisher's Exact Test. The horizontal dashed line corresponds to a nominal 5% significance-threshold with Bonferroni-correction for 250,000 tests. **b**, Magnification of the genomic region surrounding *RPM1*, the position (and extent) of which is indicated by the vertical blue line.

Figure 3 | Candidate SNPs are over-represented among strong associations. GWA analysis of the FT10 phenotype: negative log p-values from the Wilcoxon test are plotted against those from EMMA. Points corresponding to SNPs within 20 kb of a candidate gene are shown in red; the rest are shown in blue. The enrichment of the former over the latter in different parts of the distribution is shown.

Figure 4 | Association with *FLC* expression at the top of chromosome 4 near *FRI*. The p-values are from EMMA; the position of *FRI* is indicated by a vertical yellow line. **a**, Single-SNP tests. **b**, Col-allele of *FRI* (blue dot) is added as co-factor in the model. **c**, Ler-allele of *FRI* (red dot) is added as co-factor in the model. **d**, Both alleles added as co-factors in the model.